

FLiMS: A Fast Lightweight 2-Way Merger for Sorting

Philippos Papaphilippou[✉], Wayne Luk, *Fellow, IEEE*, and Chris Brooks

Abstract—In this paper, we present FLiMS, a highly-efficient and simple parallel algorithm for merging two sorted lists residing in banked and/or wide memory. On FPGAs, its implementation uses fewer hardware resources than the state-of-the-art alternatives, due to the reduced number of comparators and elimination of redundant logic found on prior attempts. In combination with the distributed nature of the selector stage, a higher performance is achieved for the same amount of parallelism or higher. This is useful in many applications such as in parallel merge trees to achieve high-throughput sorting, where the resource utilisation of the merger is critical for building large trees and internalising the workload for fast computation. Also presented are efficient variations of FLiMS for optimizing throughput for skewed datasets, achieving stable sorting or using fewer dequeue signals. Additionally, FLiMS is shown to perform well as conventional software on modern CPUs supporting single-instruction multiple-data (SIMD) instructions, surpassing the performance of some standard libraries for sorting.

Index Terms—FPGA, high-throughput, parallel merger, merge tree, sorting algorithms, sorting networks, SIMD, databases

1 INTRODUCTION

THE merge operation is widely used for a variety of applications, including in many popular sorting algorithms, such as mergesort and timsort [1], as well as in hardware for database operations, including sort-merge joins [2]. It is thus desirable to accelerate, and recent research has proposed a variety of merge accelerators on field-programmable gate arrays (FPGAs).

As frequency scaling has stopped being the primary method for achieving performance, the main way of achieving high-throughput/high-bandwidth in modern systems is now to increase the datapath width. This has influenced computer architecture in many aspects, such as with wider single-instruction multiple-data (SIMD) instructions on general purpose processors (CPUs). One application that can benefit from high-throughput is sorting, as more data are being able to be processed per cycle. This paper presents a high-throughput merger algorithm, that merges two arbitrarily long input lists with high-throughput, exporting w elements per cycle, assuming the inputs are appropriately provided, such as through banked memories, reading up to w elements from each of the two input lists per cycle.

The challenges for the merge operations on FPGAs have been the low clock frequency due to the feedback datapath

being the critical path, and the high resource utilisation in some attempts to remove the feedback datapath.

At the time of writing, FLiMS uses the least number of comparators and pipeline stages among the latest mergers. It uses a modified version of the bitonic merge block, as found in bitonic sorters (see Fig. 3), repurposed for performing 2-way parallel merge for streaming data. All alternative designs require the two input sequences of the bitonic (partial) merger (or the odd-even merge-based equivalent) to be sorted. The *main idea* is to relax this condition. This idea eliminates the need for expensive rotations before the inputs of the merger [3], or alternative workarounds involving redundant logic [4], [5], [6].

Other novel aspects in this work include the distributed nature of the selector stage, which has better timing characteristics on FPGAs, without occupying additional pipeline stages from the merger logic. The optimisation for skewed datasets is more lightweight and scalable than a previous attempt [3], as it does not rely on barrel shifters. FLiMS also does not suffer from the tie-record challenge found in all other feedback-less designs [4], [5], [6], [7], and a costly workaround is not deemed necessary. The regularity in the topology of the comparators is also found to help with SIMD implementations of high-throughput merging in CPUs.

The main contribution is the highly-efficient design of a high-throughput 2-way merger. Three variations of FLiMS are presented to achieve stable merging, high-throughput on skewed datasets and efficient memory use for special cases. Other contributions include a systematic comparison with a variety of alternative approaches, and an evaluation on both FPGAs and modern CPUs with SIMD. Automated generator scripts¹ provide (a) the Verilog code for the FPGA implementation, as well as (b) C++ code with SIMD intrinsics for CPUs, for a user-specified degree of parallelism (w).

- Philippos Papaphilippou and Wayne Luk are with the Department of Computing, Imperial College London, SW7 2BX London, U.K. E-mail: p.p.cy@ieee.org, w.luk@imperial.ac.uk.
- Chris Brooks is with the Science Innovation, dunnhumby, W6 7NL London, U.K. E-mail: Chris.Brooks@dunnhumby.com.

Manuscript received 19 June 2021; revised 15 Nov. 2021; accepted 22 Dec. 2021. Date of publication 27 Jan. 2022; date of current version 11 Nov. 2022.

This work was supported in part by dunnhumby, in part by Microsoft and in part by the U.K. EPSRC under Grants EP/L016796/1, EP/I012036/1, and EP/N031768/1.

(Corresponding author: Philippos Papaphilippou.)

Recommended for acceptance by P. Milder.

Digital Object Identifier no. 10.1109/TC.2022.3146509

1. *Source available:* <http://philippos.info/flimsj>




Fig. 1. Parallel merge tree (PMT [3]), for 8 input lists.

This paper is an extension of “FLiMS: Fast Lightweight Merge Sorter” [8] and is also partly based on the paper “An Adaptable High-Throughput FPGA Merge Sorter for Accelerating Database Analytics” [9]. It extends the material from these two papers with more detailed proofs (Section 5) and a thorough comparison with related work, such as providing exact values for the number of comparators each approach requires (Section 6). Additional variations are presented to also support stable sorting and dequeuing whole rows (Section 4). The experimental comparison is updated with implementations of the state-of-the-art alternatives [6], [7] on a high-end FPGA (Section 7), and the SIMD-based version of FLiMS is extended to implement full sorting on modern CPUs to compare with sorting library functions (Section 8).

2 BACKGROUND

On FPGAs, the available sorting accelerators are inspired by a variety of different serial algorithms, including insertion sort [9], [10], [11]. However, this background section is restricted to high-throughput merge sorting on FPGAs. Merge sorting has been one of the most studied algorithms for sorting on FPGAs. This is due to its versatility, such as when reusing the same circuits to sort arbitrarily long input recursively.

2.1 High-Throughput Merge Sorters

High-throughput merge sorters can merge a number of sorted lists simultaneously, while providing an output rate of more than 1 element per cycle. This can be achieved by building a merge tree (PMT [3]), mainly consisting of high-throughput mergers of 2 lists and FIFO queues, as with predecessors [12], [13].

Fig. 1 shows how these building blocks can be arranged to merge 8 sorted inputs of throughput 1, with an output rate 8. The ‘merge rate’ of the mergers in each level of the merge tree denotes the number of elements in their two inputs collectively and the number of elements in their output. This directly contributes to the throughput of the merger, as well as the bandwidth requirements of the proceeding and succeeding logic in the pipeline. For instance, a merger of rate 4:2 merges two inputs of width 2 (times the element width) and outputs two elements per cycle. The difference in widths from level to level is managed by rate converters and the appropriate stall signals.

One challenge in parallel merge trees is that, while they can easily saturate the available bandwidth by scaling the number of inputs, they do not scale well for high number of inputs. For this reason, many-leaf mergers have emerged, to




Fig. 2. HPMT for high-throughput and many-leaf merging.

support a higher number of inputs more efficiently (currently up to a few thousands [14], [15] rather than in tens for PMTs). Merging many inputs simultaneously helps reduce the number of data passes required for complete sorting. A single data pass is equivalent to reading the entire input data once. However, many-leaf mergers are single-rate, meaning that they can only produce one output per cycle. If the data are not wide enough, this can lead to underutilisation of the available bandwidth.

To solve these trade-offs, the hybrid parallel merge tree (HPMT) was introduced [9], to enable both high-throughput and many-leaf merging at the same time. The size of the HPMT can be easily adjusted to saturate the bandwidth of the target architecture, while eliminating the number of passes of the data by still using many-leaf merging. Fig. 2 shows how HPMT can combine 4 many-leaf mergers of K inputs (totalling $4K$ input lists) into a merge tree with an output rate of 4 elements per cycle.

2.2 High-Throughput 2-Way Mergers

The high-throughput 2-way mergers are the main building blocks of the aforementioned merge trees.

A merger for 2 already-sorted sublists of fixed length can be modified to merge 2 lists of arbitrary length in streaming fashion. Then it can be used as a building block for a parallel merge tree to merge many lists concurrently.

Most of them are based on the two popular sorting networks: Batcher’s odd-even mergesort and the bitonic sorter [16]. These two sorting networks have the same number of stages and can be built hierarchically using 2 sorters of half the input and an appropriately sized merger to merge two equally-sized sorted sublists. The merger part consumes the last $\log_2(n)$ stages in both sorting networks, where n is the number of inputs. This merge block is optimised and/or combined with additional logic to work as a parallel merger for longer lists as streams.

Fig. 3 shows the bitonic merger, as found in a bitonic sorter of 8 inputs. The pairs of circles are compare-and-swap (CAS) units, working as sorters of two inputs, i.e., $a, b \rightarrow \min(a, b), \max(a, b)$. The list $[a_0, a_1, \dots, a_7]$ is initially unsorted, and right before the merger, it is partially sorted into two sorted sublists ($[b_0, b_1, b_2, b_3]$ and $[c_0, c_1, c_2, c_3]$). The merger merges these two sublists, consuming the last 3 stages of this bitonic sorter.

The first known high-throughput merger for arbitrarily long input was based on a well known algorithm for merging using SIMD instructions on Intel processors [17]. It was adopted for database use on FPGAs [12]. However, the long feedback problem was more prominent on FPGAs [4], since




Fig. 3. A 8-to-4 bitonic partial merger as found in the bitonic sorter.




Fig. 4. Merger used in Casper *et al.* [12], [17].

it can negatively impact the critical path, and is not scalable for many inputs. The algorithm goes as follows: starting with the first w -sized batches of each of the sorted sublists, the merger produces the top w result as output to progressively merge the entire input, while the lower w result is fed back into the lower w of the input to continue the merging. A single comparison between the first element of each batch is enough to distinguish the next batch to dequeue and place at the merger. Fig. 4 shows the high-level representation of this approach for FPGAs.

Some 2-way hardware mergers (including FLiMS) utilise an optimization of the bitonic merger, the $2w$ -to- w bitonic partial merger [18], which outputs just the top half of the result (or lower if the CAS units are flipped). The bitonic partial merger is a subset of the bitonic merger. These are summarised in Fig. 3. This “pruned” merge block is combined with additional logic to work as a parallel merger for longer lists as streams.

In 2016, Song *et al.* built the parallel merge tree (PMT [3]) with $2w$ -to- w bitonic partial mergers. In Fig. 5, we can see a high-level view of the merge block. This merger works as follows: two input queues, A and B, output 0 to w elements each per cycle, according to how many made it in the last result of w elements. This is known from just the first stage of the $2w$ -to- w bitonic partial merger and is used as a feedback to select the amount of elements to be dequeued from A and B. Since some elements remain from the previous cycle, each input of the bitonic partial merger block needs to be rotated by an offset equal to the number of dequeued elements (so far). This is done to ensure that the bitonic partial merger gets sorted inputs. However, the crossbars implementing the barrel shifters create a critical path that increases with w , leading to low frequency designs and it does not scale well [19].




Fig. 5. High-level view of the merger used in PMT [3].




Fig. 6. High-level view of the mergers used in MMS [4] and VMS [5].

At some point, most of the attention was drawn on removing the expensive feedback length that existed in traditional merger designs [3], [12], that prevented scalability in terms of operating frequency for an increased degree of parallelism (w) (SHMS [19], MMS [4], VMS [5]).

In 2017, Mashimo *et al.* [19] proposed a lower-latency-feedback architecture, SHMS, to solve the long critical path problem with the previous approaches. While achieving much higher frequencies, as high as 3.14 times more than PMT [3] for 32 different input queues (also focusing on multiple inputs), the register utilisation was much higher (7.26 times more registers than PMT for 32 inputs). This does not scale well, not from the long feedback datapath length, but due to the high register utilisation.

Then, Saitoh *et al.* [4] proposed a feedback-less architecture, MMS, to increase the performance and scalability of the merge operator. MMS uses two $2w$ -to- w bitonic partial merge blocks along with shift registers and an extra comparator and multiplexer. VMS [5] is a variation of MMS that is based on odd-even mergers instead, but also focuses on improving the tie-record workaround (see Section 6). In Fig. 6, we can see the high-level view of these designs.

Finally, FLiMS [8] and WMS/EHMS [6] offered further improvements by focusing on efficiency, for minimising the required hardware resources, usually with a subsequent improvement in operating frequency. Essentially, WMS is an optimisation of MMS [4] (and VMS [5]), because it fuses the two $2w$ -to- w partial mergers into one bigger merger block, and also eliminates the need for additional shift registers. Fig. 7 shows how a single merger is used in WMS, closely resembling the other approaches in MMS, VMS and EHMS.

EHMS [7], was then introduced as a potential successor to EHMS [6]. EHMS and EHMS try to move some complexity to the selector stage for lower resource utilisation at the expense of a lower operating frequency. EHMS specifically is not considered here for comparison, as its resource utilisation is close to EHMS, but with a further performance




Fig. 7. Dequeue architecture used in WMS [6].

overhead due to the increased complexity of its selector stage, worsening its critical path [7]. Also, in contrast to the other high-throughput mergers, EHMSP is for w values not in powers of 2, making it less versatile.

3 A NOVEL 2-WAY HIGH-THROUGHPUT MERGER

FLiMS is a novel 2-way high-throughput merger that only uses a single $2w$ -to- w bitonic partial merger. It merges 2 sorted inputs with high throughput without the use of barrel shifters or shift registers.

In Fig. 8, we can see a high-level visualisation of the proposed parallel merger. In contrast to previous works, FLiMS is shown to dequeue from the banks on an individual basis, rather than in batches of w elements, due to the decentralised selector logic.

In Fig. 9, we see a lower-level representation, where the green circle pairs are the compare-and-swap units, with the exception of the first pipeline stage, which has w pairs of one green and one grey circle, representing the MAX units (selector stage). If we ignore the modification of the first stage to MAX units, this topology is the bitonic partial merger [18], and it would produce the top w out of two sorted lists of w elements, and is a subset of the bitonic merger (see Fig. 3).

The proposed parallel algorithm can be broken down into 3 segments: a *selector stage* for handling the input from multiple banks/queues, the pipelined *bitonic partial merger* (minus the first stage) and finally the *output logic*.

3.1 Selector Stage

After the data are written into the BRAM banks (stored with a round-robin priority), a set of independent entities (nodes) are responsible for controlling the input to the merger. These w entities respond to the same clock and their behaviour can be described by a distributed algorithm.

Let A and B be the input FIFO queues containing the two sorted lists to be merged. Let $A_0, \dots, A_{w-1}$ and $B_0, \dots, B_{w-1}$ be the queues corresponding to the respective banks for A and B . Each of the entities $MAX_0, MAX_1, \dots, MAX_{w-1}$ have as input the pair of queue heads $(a_0, b_{w-1}), (a_1, b_{w-2}), \dots, (a_{w-1}, b_0)$ respectively, with a_i being the head of A_i and b_i being the head of B_i . Each of these entities outputs one number per cycle when both inputs are valid.





Fig. 8. Dequeue architecture of FLiMS [8].

Fig. 9. FLiMS in low-level: MAX selector stage plus a CAS network (butterfly topology), $w = 4$.

Each entity MAX_i has the data registers cA_i and cB_i to store the last heads that were dequeued from banks A_i and B_{w-1-i} respectively. It also has a register for the sorting network input (in_i). On each cycle with valid input, if $cA_i > cB_i$, it means that cA_i will make it into the top w in the result of the pipeline and therefore cA_i is copied into in_i . In this case, cA_i is replaced by the head a_i , which is dequeued from A_i , but cB_i will remain unchanged, since it will need to be compared again in the next cycle, being in the lower w . The equivalent logic goes for the case when $cA_i \leq cB_i$.

In algorithm 1, we can see the pseudocode for the distributed algorithm. Collectively, this algorithm replaces the first stage of the partial bitonic merge (half-cleaner), with MAX units instead of CAS units. It selects the current top w on each cycle and inserts them into pipeline registers for the rest of the CAS network to sort and produce the correct w -sized chunk of output.

Algorithm 1. MAX_i Unit pseudocode

```

1: int  $i$                                 ▷  $i$  is the entity tag
2: reg  $cA_i, cB_i, in_i$                  ▷ registers of data width
3: while forever do
4:   receive (positive clock edge);
5:   if  $cA_i > cB_i$  then
6:      $in_i \leftarrow cA_i$ ;
7:      $cA_i \leftarrow dequeue(a_i)$ ;
8:   else
9:      $in_i \leftarrow cB_i$ ;
10:     $cB_i \leftarrow dequeue(b_{w-1-i})$ ;
11:  end
12: end

```

There may be a need for some extra logic required to correctly handle the ending of the input queues, but it is omitted here for the sake of simplicity and portability to different architectures, as it is relatively trivial to construct. For example, when sorting natural numbers in descending order, the value 0 can be passed afterwards to handle the ending without additional dedicated logic.

3.2 CAS Network

The compare-and-swap (CAS) network of FLiMS is responsible for sorting the top- w result of the selector stage. It is a partial bitonic merger minus the first stage, or in other words, a butterfly network. It is not considered a sorting network on its own and does not sort arbitrary number sequences. For its input though, it behaves correctly and its output is always sorted (see proof 5.1).

TABLE 1
Merging Two Descending Sequences: Example Execution for $w=4$ and Random Sorted Lists A and B¹

Cycle	Input A	Input B	cA	cB	Output (after the pipeline delay)
0	3 3 4 5 11 16 17 26 26 29	0 7 8 9 12 15 18 19 21 22			
1	3 3 4 5 11 16	0 7 8 9 12 15	17 26 26 29	22 21 19 18	22 26 26 29
2	3 3 4	0 7 8 9 12	17 5 11 16	15 21 19 18	17 18 19 21 22 26 26 29
3	3 3	0 7	4 5 11 16	15 12 9 8	11 12 15 16 17 18 19 21 22 26 26 29
4			4 5 3 3	7 0 9 8	5 7 8 9 11 12 15 16 17 18 19 21 22 26 26 29
5			4 3 3	0	0 3 3 4 5 7 8 9 11 12 15 16 17 18 19 21 22 26 26 29

¹see http://philippos.info/sort_visual for an online visualisation.

3.3 Output Logic

On each cycle, if the output of the partial bitonic merger is marked as valid, it is written down as a w -sized chunk of the result, such as in w output banks containing implementing an output queue O .

One observation is that when there is valid output per cycle, it produces exactly w elements, as with other 2-way mergers. This is useful for easing synchronisation when embedding into a merge tree [9].

The sets of registers cA and cB are not visible in the first (selector) pipeline stage of Fig. 9. The notion of those registers is optional because they can also be considered the current sets of heads of the banked queues A and B. Though, it is sometimes convenient to use cA and cB , such as when the input queues are block RAM sections, where a read register is already present for reading each memory.

Table 1 presents an example execution for $w=4$.

4 ADDITIONAL FUNCTIONALITY

This section presents variations of FLiMS, which can be used to increase its applicability or performance, according to the distribution of the data and the requirements of the sorting problem and platform.

4.1 Skewness Optimisation

FLiMS can be used to build parallel merge trees that merge many input lists hierarchically in a single pass. Parallel merge trees can suffer from rate mismatch that occurs when the input data distribution leads to underutilisation of certain mergers, resulting in reduced throughput [3], [14], [15].

While the memory access throughput is a matter of the memory system, it might be allowed for the accelerator to receive the data from each of the input lists A and B with a fixed bandwidth, less than w , such as in a PMT. A contributor to rate mismatch is when there are a lot of duplicates in the input (skewed datasets). When the data are skewed, the merger only dequeues from one of the input queues for long periods of time. This results in stalled cycles from underutilising the aggregate bandwidth of the queues.

PMT [3] proposes a simple solution which causes the merger blocks to fetch from both inputs at a similar rate when there are duplicates. This ensures that the input queues are consumed with a similar throughput, that collectively balances the utilisation of the merge tree. However, PMT's mergers inherit the long feedback problem, which was addressed in subsequent works [4], [5], [6], [8].

We propose the equivalent optimisation for FLiMS [8], while keeping the decentralised nature of the selector stage.

On duplicates, there is an “oscillating” effect at the MAX units, which balances the dequeuing rate from the two groups of inputs. The code for the new selector units is illustrated in algorithm 2. An 1-bit register called dir_i represents the input out of which the result was taken during the previous cycle, and is appended to the least significant bit in the comparison, to enforce a sort priority on equal values.

Algorithm 2. Modified MAX_i Unit pseudocode for the Skewness Optimisation

```

1: int  $i$ ; ▷  $i$  is the entity tag
2: reg  $cA_i, cB_i, in_i$ ; ▷ registers of data width
3: reg  $dir_i$ ; ▷ 1-bit register
4: while forever do
5:   receive (positive clock edge);
6:   if  $\{cA_i, dir_i\} > \{cB_i, !dir_i\}$  then
7:      $in_i \leftarrow cA_i$ ;
8:      $cA_i \leftarrow \text{dequeue}(A_i)$ ;
9:      $dir_i \leftarrow 0$ ;
10:  else
11:     $in_i \leftarrow cB_i$ ;
12:     $cB_i \leftarrow \text{dequeue}(B_{w-1-i})$ ;
13:     $dir_i \leftarrow 1$ ;
14:  end
15: end

```

4.2 Stable Merge

In contrast to the skewness optimisation, stable sort may be desired instead. Stable sort is when the sorted output has the same order for duplicate values as they appear in the input. For implementing a stable merge sort, FLiMS would also need to be stable, i.e., to prioritise the duplicates of input A over the ones from input B, and keep their original order inside A and B accordingly. Such a modification cannot co-exist with the skewness optimisation, since the priority between the duplicates will be based on the input source.

Originally, FLiMS is not stable, as it is partly-based on the bitonic sorter, which is not stable. Temporarily appending the input source (1 bit) and the port number ($\log_2(w)$ bits) to the MSB would be required to disambiguate between the original order of duplicates inside the CAS network. As the order of the inputs inside MAX units are naturally rotated by an offset, the port order is not enough to distinguish the order inside each batch containing duplicates. For this reason, a 2-bit value needs to be carried between the input source and the port number, that keeps count of the batch order. A single-bit counter would not be enough for distinguishing which of the two compared entries came first.




Fig. 10. FLiMSj merging example for dequeuing whole rows, $w = 4$.

Algorithm 3 shows the modifications required on the MAX_i units to support stable merge.

Algorithm 3. Modified MAX_i Unit pseudocode for Implementing Stable Merge in Descending Order

```

1: int  $i$ ; ▷  $i$  is the entity tag,  $\log_2(w)$ -bits
2: reg  $cA_i, cB_i, in_i$ ; ▷ registers of data width
3: reg  $orderA_i = 0, orderB_i = 0$ ; ▷ 2-bit registers
4: while forever do
5:   receive (positive clock edge);
6:   if  $cA_i > cB_i \parallel cA_i == cB_i$  then
7:      $in_i \leftarrow \{1, orderA_i, w - 1 - i, cA_i\}$ ;
8:      $cA_i \leftarrow \text{dequeue}(A_i)$ ;
9:      $orderA_i \leftarrow orderA_i - 1$ ;
10:  else
11:     $in_i \leftarrow \{0, orderB_i, i, cB_i\}$ ;
12:     $cB_i \leftarrow \text{dequeue}(B_{w-1-i})$ ;
13:     $orderB_i \leftarrow orderB_i - 1$ ;
14:  end
15: end

```

Additionally, the CAS units also need to be modified to correctly prioritise the case where the 2-bit order is “00” against “11”, as all other combinations (same values or other pairs having a difference of one) would correctly represent the original order priorities. The general idea of this approach is to emulate appending the original input order to the MSB of the data, but with a steady and low number of bits for merging arbitrarily long input. The order field can be seen as the last few bits of the input order.

4.3 Dequeuing Whole Rows (FLiMSj)

One potential advantage of the majority of the related work [4], [5], [6], [7], [12] over FLiMS is that they dequeue whole rows of w elements (or $w/2$ for EHMS [6]) from the inputs by default (see Fig. 7). This reduces the number of dequeue signals, and can also be more efficient in special cases, such as when reading narrow data from wider memories.

A relatively efficient buffering arrangement can unify the dequeue signals for FLiMS as well. This is possible because the FIFOs in FLiMS are collectively consumed in round-robin fashion and at no point two FIFO indexes of the same input differ by more than one. This means that a set of registers for buffering the next queue heads is enough for dequeuing an element batch at the right time, while providing full bandwidth for the respective input. This can be achieved with a single set of w registers for both inputs.

Fig. 10 introduces the related modification to the MAX units using a merging example for $w = 4$. The general idea is that the top $2w$ to w elements can be stored in a set of w registers (cR) after every selection iteration (cycle), while maintaining their natural rotation order. This results in at least w available elements per input, when combined with the current heads in cA and cB , while still eliminating the need for a rotation and a growing feedback. Algorithm 4 describes this approach in more detail.

Algorithm 4. Modified MAX_i Unit pseudocode for Dequeuing Whole Rows From the Inputs (FLiMSj)

```

1: int  $i$ ; ▷  $i$  is the entity tag
2: reg  $cA_i, cB_i, cR_i, in_i$ ; ▷ registers of data width
3: reg  $dir_i, src_i$ ; ▷ 1-bit registers
4: while forever do
5:   receive (positive clock edge);
6:   ▷ (use the correct head based on  $src_i$ )
7:   if  $(src_i ? cA_i : cR_i) > (src_i ? cR_i : cB_i)$  then
8:      $in_i \leftarrow (src_i ? cA_i : cR_i)$ ;
9:      $dir_i \leftarrow 0$ ;
10:  else
11:     $in_i \leftarrow (src_i ? cR_i : cB_i)$ ;
12:     $dir_i \leftarrow 1$ ;
13:  end
14:  sync( $dir_i$ ); ▷ (data dependency)
15:  if  $src_i == dir_i$  then
16:    ▷ (the consumed element is from  $cR_i$ )
17:     $src_i \leftarrow dir_0$ ; ▷ (next following  $MAX_0$ )
18:     $cR_i \leftarrow dir_0 ? cB_i : cA_i$ ;
19:  end
20:  ▷ Fetch the next batch collectively
21:   $dir_0 ? cB_i \leftarrow \text{dequeue}(b_{w-1-i}) : cA_i \leftarrow \text{dequeue}(a_i)$ ;
22: end

```

5 CORRECTNESS

This section provides proofs on some non-trivial parts of our proposed techniques. The correct operation of FLiMS and its skewness optimisation can be proven by induction. FLiMS can be used for merging lists in descending order, as well as for merging ascending lists with minor modifications (reversing all comparators, and reversing the order in the stable version). Here, we study FLiMS designs with data in descending order, but the proofs can be easily adapted for the other case, without loss of generality.

5.1 Without Additional Functionality

In order to prove that the main design behaves as expected we will show that it is functionally equivalent to a more trivial merger implementation, where on each cycle the input comes sorted to a $2w$ -to- w bitonic partial merger, as with the merger used in PMT [3]. We will show that (1) the selector stage algorithm always selects the top w out of the $2w$ -sized input and that (2) the butterfly network always sorts this top w list before writing the result in output.

(1) We denote as $l_A, l_B \in \{0, 1, \dots, w-1\}$ the naturally-occurring rotation offsets for input queues A and B respectively. Supposedly, the selector stage dequeues consecutive

elements from each banked input. Since A and B are written in a round robin fashion inside the banked memory, different rotation offsets would be required to read a sorted set of w elements from each set of FIFOs, similar to a design that rotates its inputs like PMT [3]. Note that the *MAX* units in FLiMS receive elements from the corresponding banks without performing any additional rotation.

Induction hypothesis: On each cycle, if the selector stage has worked correctly on the previous cycle, it will load $k \in \{0, 1, \dots, w\}$ elements from A and $w - k$ from B, collectively corresponding to the combined top w out of the available $2w$ elements.

Base step: At cycle 0, the parallel merger behaves in the same way as the $2w$ -to- w bitonic partial merger, because the w inputs from each list are already sorted ($l_A = 0$ and $l_B = 0$), corresponding to the first w elements from A and first w from B (stored in cA and cB registers respectively). The first stage of the $2w$ -to- w bitonic partial merger is known as a half-cleaner, and produces the top w out of the $2w$ inputs, in a bitonic sequence.

Induction step: Given that the selector stage worked correctly in the previous cycle, $k' \in \{0, 1, \dots, w\}$ and $w - k'$ are the amounts of sorted elements dequeued last from A and B respectively.

On each cycle, l_A and l_B will be updated according to the number of dequeued heads, as the starting positions shall succeed the dequeued elements. That is, $l_A = (l'_A + k') \bmod w$ for A, and $l_B = (l'_B + (w - k')) \bmod w$ for B, where l'_A, l'_B the offsets of the previous cycle. Therefore,

$$\begin{aligned} l_B &= (l'_B + (w - k')) \bmod w \\ \Rightarrow l_B &= (l'_B + (w - (l_A - l'_A))) \bmod w \\ \Rightarrow l_B &= (l'_B - l_A + l'_A) \bmod w \\ \Rightarrow (l_B + l_A) \bmod w &= (l'_B + l'_A) \bmod w \end{aligned}$$

As cycle 0 assumes $l'_A = 0$ and $l'_B = 0$ (due to the input already being aligned correctly in the banks), and that the induction hypothesis is assumed correct for all previous cycles, it always holds that $(l'_B + l'_A) \bmod w = 0$. Since $l'_B = l'_B \bmod w$, $l_B = (w - l_A) \bmod w$ for every cycle.

In order for the selector stage to produce the top w out of the current $2w$ elements, the comparisons that need to be made are between all pairs in $\{(Ta_i, Tb_{w-1-i}) | \forall i \in \{0, 1, \dots, w-1\}\}$, where $Ta_0, \dots, Ta_{w-1}$ are the current top w elements in A and $Tb_0, \dots, Tb_{w-1}$ are the current top w elements in B, sorted in descending order. These comparisons are required in order to emulate the first stage (half-cleaner) of the bitonic partial merger.

The *MAX* units have the same topology as the bitonic (partial) merger (as seen in Fig. 3). We notice that whatever the rotation combination from l_A and l_B is, the correct comparisons will be made, as

$$\begin{aligned} &(Ta_{(i+l_A) \bmod w}, Tb_{((w-1-i)+l_B) \bmod w}) \\ \equiv &(Ta_{(i+l_A) \bmod w}, Tb_{((w-1-i)+(w-l_A)) \bmod w}) \\ \equiv &(Ta_{(i+l_A) \bmod w}, Tb_{(w-1-i-l_A) \bmod w}) \\ \equiv &(Ta_{(i+l_A) \bmod w}, Tb_{(w-1-(i+l_A)) \bmod w}). \end{aligned}$$

(2) The sorting network receives the correct top- w output from (1) and the task is to sort it. Originally, this structure is supposed to sort bitonic sequences of size w [16]. The output

of the distributed algorithm block is a rotated bitonic sequence, as we saw that the comparisons will be rotated by l_A . A rotated bitonic sequence is also a bitonic sequence [20], therefore the input for the sorting network has the correct property. This completes the proof.

5.2 Including the Skewness Optimisation

To prove that FLiMS continues to sort correctly, the selector stage must be shown to still produce a bitonic sequence [8] (up to one local maximum and up to one local minimum).

The bitonic sequence property needs to apply also on the order of each input when there are duplicates. Therefore, we need to show that there will still be up to one local maximum and up to one local minimum in the bitonic sequence, even though there might be multiple additional entries with a value equal to the minimum or maximum.

The original order between consecutive duplicates in the same input is used to correctly prioritise duplicates, as the input lists are considered already sorted. Being consistent about the original order ensures dequeuing consecutive entries from A and B, keeping the integrity of the input data, as the data are stored in round-robin fashion in banks.

On each cycle, each *MAX_i* units compares Ta_j to Tb_{w-1-j} , where i is a rotation of j by a common offset $o \in \{0, 1, \dots, w-1\}$ ($o = l_A = (w - l_B) \bmod w$), and Ta and Tb represent the remaining elements of the input queues in descending order. This emulates a half-cleaner that selects (and dequeues) a total of the greatest w elements from a total of $2w$ elements, generating a bitonic sequence.

Induction hypothesis: On each cycle, if the skewness-optimised selector stage worked correctly on the previous cycle, it will produce a bitonic sequence from k consecutive elements from A and $w - k$ consecutive elements from B, where $k \in \{0, 1, \dots, w\}$.

Base step: At cycle 0, the common offset is zero ($o = 0$), as the inputs are properly aligned inside the input FIFOs, and no element has been dequeued yet. Each *MAX_i* unit has an initial value of 0 stored in its dir_i register. This means that on the event of comparing duplicates, it will behave as in the non-optimised case, where a single source (B) is preferred for duplicates, as the comparison of line 6 of algorithm 2 is now equivalent to line 5 of algorithm 1 ("if $cA_i > cB_i$ then"), which is already proven to produce a bitonic sequence from the proof of Section 5.1.

Induction step: The skewness optimisation modification only takes effect where there are duplicates, i.e., $\exists h \in \{0, 1, \dots, w-1\} : Ta_h = Tb_{w-1-h}$. In such a case, we notice that this only happens consecutively (including wrap-arounds from the natural rotation) and for the minimum value of the output, as $(Ta_0, Ta_1, \dots, Ta_{w-1})$ is monotonically decreasing and $(Tb_{w-1}, Tb_{w-2}, \dots, Tb_0)$ is monotonically increasing. As a consequence, the position of the minimum (split) in the bitonic sequence can be at the start, end or between this region of duplicates.

Given that FLiMS worked correctly on the last cycle, the dir_i registers correspond to the last half-cleaner decisions, which will be of the form $\{1\}^{w-k}\{0\}^k$, $k \in \{0, 1, \dots, w\}$, after considering the offset o . Ones and zeros appear consecutively, since only consecutive elements are dequeued from each input list on each cycle. Also, the sequence of dir_i registers starts from 1, when we consider the current rotation offset.

TABLE 2
Comparing High-Throughput 2-Way Mergers

2-way merger	Feedback length	Latency	Number of comparators	H/W modules	Merger topology	Tie-record challenge
basic [12], [17]	$\log_2(w) + 2$	$\log_2(w) + 2$	$w + w\log_2(w)$	$1 \times 2w\text{-to-}2w$ merger	bitonic	no
from PMT [3]	$\log_2(w) + 1$	$2\log_2(w) + 1$	$w + \frac{1}{2}w\log_2(w)$	$1 \times 2w\text{-to-}w$ merger & 2 barrel shifters	bitonic	no
MMS [4]	1	$2\log_2(w) + 3$	$2w + w\log_2(w) + 1$	$2 \times 2w\text{-to-}w$ mergers & shift registers	bitonic	yes
VMS [5]	1	$2\log_2(w) + 3$	$2w + w\log_2(w) + 1$	$2 \times 2w\text{-to-}w$ mergers & shift registers	odd-even	yes
WMS [6], [7]	1	$\log_2(w) + 3$	$3w + \frac{1}{2}w\log_2(w)$	$1 \times 3w\text{-to-}w$ merger	odd-even	yes
EHMS [6], [7]	1	$\log_2(w) + 3$	$\frac{5}{2}w + \frac{1}{2}w\log_2(w) + 2$	$1 \times 2.5w\text{-to-}w$ merger	odd-even	yes
FLiMS [8]	1	$\log_2(w) + 1$	$w + \frac{1}{2}w\log_2(w)$	$1 \times 2w\text{-to-}w$ merger	bitonic	no
FLiMSj	1	$\log_2(w) + 2$	$w + \frac{1}{2}w\log_2(w)$	$1 \times 2w\text{-to-}w$ merger	bitonic	no

This is because the naturally-occurring rotation offset is updated according to the last position dequeued from A , the next of which corresponds to the first from list B , that yields the “first” 1 in the dir_i sequence.

Therefore, the region of duplicates will be a sublist of the expression $\{1\}^{w-k}\{0\}^k$, with its 1s and 0s replaced by consecutive duplicates from A and B respectively. As a result, there will be up to one local minimum (split) in this region, and therefore up to one local minimum in the entire half-cleaner result, which consists a bitonic sequence.

6 COMPARISON WITH THE RELATED WORK

Table 2 compares FLiMS to the related work, according to different terms contributing to the resource utilisation and efficiency of the design. FLiMS uses the least amount of resources by only requiring a single $2w\text{-to-}w$ bitonic partial merger. It has a single-stage feedback latency, making it “feedback-less”, and has the least amount of latency, which is $\log_2(2w)$. FLiMSj of Section 4.3 is also added to the table, even though its only modification is in its MAX units, which results in one more cycle of pipeline latency.

The first two entries [3], [12] have a feedback consuming multiple stages, such as from the additional $\log_2(w)$ stages required to implement the barrel shifters before the inputs in the mergers of PMT [3]. Hence, an increased number of inputs has a scalability problem, as a pipelined implementation of the feedback would reduce the throughput. Alternatively, squeezing the increasing logic into a single pipeline stage (consuming a single cycle) can heavily impact the operating frequency [4].

MMS [4] and VMS [5] were the first solutions that provided a practical solution for the feedback problem with a relatively low resource utilisation. Their approach was to use either two $2w\text{-to-}w$ bitonic partial mergers (MMS) or two $2w\text{-to-}w$ odd-even mergers (VMS). Both of those topologies have $\log_2(2w)$ pipeline stages and are relatively similar. They are from the last $\log_2$ steps from the bitonic sorter and odd-even merge sort respectively [16].

WMS and EHMS [6] on the other hand, achieve to use a single feedback-less merger, as with FLiMS, but this merger is for double the inputs (plus optimisations), totalling to one more pipeline stage.

Fig. 11 shows how a $4w\text{-to-}4w$ merger from odd-even mergesort is adopted to implement $3w\text{-to-}3w$, by pruning CAS units. There is, though, additional pruning as the output is only w elements. EHMS uses the same merger, but the first $w/2$ values of the input are not used, resulting in

fewer comparisons. The resulting amount of comparators for each approach is shown in the table 2, and the formulas mainly derive from Cullen numbers [21], and are validated by using yosys through synthesising the Verilog implementations of the evaluation Section 7.

The pipeline length (latency) also impacts the resource utilisation, as certain values need to be propagated for longer through the pipeline registers. The merger of WMS and EHMS uses one more cycle than FLiMS, as it is an optimised merge block for double the inputs (as seen in Figs. 9 and 11). Note that WMS and EHMS propose an optimisation to reduce the pipeline registers by fusing some CAS units, but it is ignored in this comparison, as it can be explored separately for all mergers. All mergers other than FLiMS and in PMT [3], have a separate single-cycle selector stage, contributing to one more pipeline stage, while in FLiMS it is integrated in the modified bitonic partial merger.

One challenge with the mergers MMS [4], VMS [5], WMS and EHMS [6] is that, if there are duplicate values being compared, the output can be corrupted, also known as the *tie-record issue*. Specifically, this is a problem in key-value pairs, where only the key is compared, and the integrity of the values can be lost where there are duplicate keys. The available workarounds of the related works vary in complexity (not presented). In FLiMS, this is not the case, as the selector stage decides for the top w result to propagate through the output immediately. In the other approaches, finding the top w is done more indirectly, relying on two orders, and the problem arises due to the odd-even merge and bitonic sort topologies not implementing stable sort.

7 FPGA IMPLEMENTATION

In order to evaluate FLiMS on FPGAs, we compare its resource utilisation and maximal operating frequency with




Fig. 11. Merger used in WMS and EHMS [6].

TABLE 3
Resource Utilisation as AXI Peripherals, as Reported by Vivado

w	FLiMS		FLiMSj		WMS		EHMS	
	kLUT	kFF	kLUT	kFF	kLUT	kFF	kLUT	kFF
4	1.7	2.9	2.5	3.2	2.7	5.3	3.1	4.8
8	3.6	6.3	5.1	6.8	5.6	11.0	6.2	10.3
16	7.0	1.4	10.6	14.6	11.7	23.1	13.0	21.6
32	15.4	29.0	20.9	31.2	23.5	48.3	26.7	45.3
64	33.7	62.0	45.0	66.4	53.3	100.8	57.9	94.6
128	73.4	132.2	96.1	140.8	106.6	209.8	120.4	197.5
256	158.6	280.7	208.6	297.9	224.0	436.0	252.2	411.4
512	345.3	594.0	436.2	628.4	473.0	904.7	525.3	855.6

the latest alternatives WMS and EHMS [6]. Our comparison includes FLiMSj for including a more direct competitor to the state-of-the-art WMS and EHMS mergers, when dequeuing whole rows is required (see Section 4.3). All generator scripts are implemented from scratch and produce Verilog code for each of the compared mergers for a given degree of parallelism w and data width. This experiment uses 64-bit mergers and targets the Xilinx Alveo U280 board.

The generated designs work as simple AXI peripherals, that read already sorted data stored in distributed memory (on-chip) and write back also to distributed memory. The host places the sorted sublists and reads the merged result for validation purposes. For every different value of w , the FIFOs are only 2 elements deep (totalling $4w$ elements for input and output) to eliminate the differences between different merger designs to only their core logic. (Though, the bitstreams generated for validation had longer queues).

In order to simplify the comparison, the fusion of some pairs of compare-and-swap units (CAS) in WMS and EHMS [6] is not followed in this evaluation. This omitted optimisation could be explored separately, as with removing pipeline registers [22], and does not directly relate to the main structure of the mergers. Additionally, the tie-record workarounds of WMS and EHMS have not been taken into consideration, even while FLiMS does not suffer from the tie-record issue. Thus, unique input values are assumed, such as with timestamp information inside the 64-bit input, or no satellite/payload data (i.e., values in key-value pairs).

Table 3 includes the dataset on the obtained resource utilisation of look-up-tables (LUTs) and flip-flop registers (FF), as reported by Vivado 2020.1. Fig. 12 is based on the same data and uses FLiMS as a baseline to emphasise on the overheads of the alternative approaches on resource utilisation. As a conclusion, FLiMS uses the least amount of LUTs and FFs, while WMS and EHMS only differ marginally, as they are based on a similar merger. As expected [7], between WMS and EHMS, WMS wins in LUT utilisation, while EHMS wins in FF utilisation. Using the current implementation as an AXI peripheral, FLiMS is roughly about 1.5 to 2 times more hardware resource efficient. FLiMSj has almost the same FF utilisation as FLiMS, though in terms of LUTs it is about 1.3x more expensive than FLiMS, but always more resource efficient than WMS and EHMS.

Finally, Fig. 13 presents the comparison of the obtained maximal operating frequencies through the reported worst negative slack (WNS). Most datapoints used the default




Fig. 12. Comparing resources of the state-of-the-art, FLiMSj and FLiMS.

Vivado 2020.1 settings, though additional directives such as aggressive explore were used on some outliers or non-routable designs, especially for $w \geq 128$. Having such irregularity or small variations in the results are expected, as place-and-route is heuristic-based and becomes more challenging for larger designs. For WMS with $w \geq 256$, the additional tested directives did not help with routability.

FLiMS has a considerable advantage over both WMS and EHMS, sometimes yielding more than double the operating frequency. FLiMSj has a small overhead over FLiMS, though WMS seems to marginally win for $w \leq 16$. WMS is known to be better performing than EHMS, at the expense of additional hardware resources, though the reduced resources help with routability and performance for high values of w . FLiMS wins in both performance and resource utilisation by a great margin, while FLiMSj lays between FLiMS and the alternatives WMS/EHMS in most aspects to offer the additional functionality of Section 4.3.

This evaluation focused more on the merging techniques, ignoring the building of merge trees or the handling of the list endings. The relevant paper [9] elaborates on a complete sorter implementation based on FLiMS, with a highly competitive logic and time complexity combination over related work on complete sorting.

8 SOFTWARE IMPLEMENTATION USING SIMD INTRINSICS

The goal of this section is to experimentally show that a single-instruction multiple data (SIMD)-accelerated merge sort




Fig. 13. Maximal operating frequency for FLiMS(j), WMS and EHMS.

function based on FLiMS can compete with alternative popular sorting functions based on different sorting algorithms. Today's general purpose processors (CPUs) feature vector or SIMD instructions as a way to increase performance in numerous compute and memory-intensive applications. Parallel merging algorithms implemented using SIMD-intrinsics, have already been shown to improve the sorting performance on CPUs.

Chhugani *et al.* [17] used the rather simple merge algorithm based on the bitonic merger (see Fig. 4) to enable high-throughput merging on an older Intel processor. Since it uses a full (non-partial) bitonic merger, both the lower w and upper w are used. As the upper w is calculated after $\log_2(2w)$ stages, this could have the feedback problem that is addressed by the latest research on FPGA merging (as summarised in table 2). However, on CPUs this is not much of a concern, because the pipelining functionality is not that advanced to achieve efficient task-pipelining. For example, for a single layer of compare-and-swap (CAS) units, there need to be at least three SIMD intrinsics, one for *min*, one for *max* and at least one *shuffle*. The latter is to appropriately permute the inputs to emulate the CAS network topology.

In practice, how FLiMS could help is with a reduction in instruction count, as the lower w of the result is not needed. In other words, with roughly a similar number of instructions, FLiMS can merge with double the amount of parallelism w . Other desirable characteristics of FLiMS on SIMD are the bitonic merge topology, which is more regular/ symmetric when compared with the odd-even merge sort topology, and the elimination of the rotation of the inputs. A similar discussion can be made for comparing with the other merger alternatives of Section 6, which require more comparisons and lengthier pipelines than FLiMS.

In order to assess the efficiency of FLiMS as an SIMD algorithm, a manually vectorised code in C++ is developed using Advanced Vector Extensions for 256-bit registers (AVX2). It is then extended to perform full sorting, and is finally compared to other existing sorting functions.

8.1 Merge Function Implementation

An SIMD-based implementation of FLiMS can be split into two main parts; one for the *MAX* (selector) stage and one for the butterfly network. The *MAX* stage is responsible for fetching the next elements from the input lists, selecting the top w and feeding it to the butterfly network. The butterfly network is more straightforward to implement using intrinsics, as it is a part of the bitonic sorter, which was already explored in a similar context [23].

The *MAX* stage can be implemented in SIMD by keeping the *cA* and *cB* in vector registers and generating the top w by comparing them. There are two ways to implement its fetching functionality. The first is to keep and update w independent pointers per input representing the input queues. Then mask/gather intrinsics are called to update the *cA* and *cB* values which are kept in vector registers. This is more faithful to the original FLiMS algorithm, but it is less efficient to use a gather AVX2 intrinsic [24] for fetching otherwise continuous memory locations. The preferred faster




Fig. 14. Finding the optimal w value for AVX2-based FLiMS.

method involves “pre”-fetching w -sized batches of elements, which is reminiscent of FLiMSj of Section 4.3.

One complication with fetching whole vectors for the *MAX* stage is that the *cB* vector needs to be in reverse order. This is done by the intrinsic `_mm256_permutevar8x32_epi32()` to reverse the contents of each batch fetched from list B. The comparison of the *MAX* units is done by the intrinsic `_mm256_cmpgt_epi32()` and the result boolean vector is also used to fetch the next heads selectively after blending `_mm256_blendv_epi8()` with the current set of heads. The result of the compare instruction `_mm256_cmpgt_epi32()` has the form of zeros and ones, each now denoting the source list of the next heads. The negation `_mm256_andnot_si256()` of this vector is also used to get the equivalent vector to be used when blending the next vector from the second input.

With respect to the butterfly network part, each layer of CAS units can be calculated by calling both a *min* (`_mm256_min_epi32()`) and a *max* (`_mm256_max_epi32()`) instruction consecutively, since each CAS has two outputs. Alternatively, a compare instruction can be used once `_mm256_cmpgt_epi32()`, but there is a performance overhead from the additional calls of the `_mm256_blendv_epi8()` instruction, required to translate the result vector of `_mm256_cmpgt_epi32()` into minimums and maximums. Before each CAS layer, the inputs must be properly aligned, and this is done through different permute intrinsics such as `_mm256_shuffle_epi32()` and `_mm256_permute2x128_si256()`.

In principle, emulating a wider FLiMS, can result in increased data locality, as more computation is done in registers. On the other hand, more logic needs to be represented, and the instruction count of the loop increases, with the additional disadvantage of less-obvious data dependencies. Since the number of vector units are limited in processors, emulating more logic on the merge loop can trigger more cache accesses and worsen the performance.

Fig. 14 shows how the degree of parallelism w of the emulated FLiMS influences the achievable throughput on an AVX2-native processor (Intel i7-8809G, with a steady 4.2 GHz clock). The C++ code for the 2-way merge function is generated by a python script, and is compiled with `-O3` and `-march=native`. Two sorted random inputs of 2^{24} elements are fed into the FLiMS merge function. The conclusion is that at $w = 16$ and 32 the throughput is the highest, and there is little variation between different compilers.

8.2 Complete Sorting

The FLiMS-based CPU merge function can be used recursively to merge input of arbitrary length, by accelerating merge sort with SIMD instructions.




Fig. 15. Evaluating the FLiMS-based SIMD complete sort C++ function.

As a complement to the FLiMS-based merge function, a sort-in-chunks function is developed to facilitate the need for initial sorted chunks, as well as to provide long-enough chunks for FLiMS to benefit from streaming access patterns. This function is based on the bitonic sorter, with a similar technique to Section 8.1 for building the butterfly network. A similar approach was followed in an SIMD-based quick-sort implementation [23], although here it is implemented from scratch with AVX2 intrinsics. In our use case, the optimal sorted chunk size is found to be 512 integers.

The performance of the FLiMS-based sort function is compared against the C++ Standard Library implementation of sorting `std::sort()`, as well as a highly-optimised [23] SIMD-based radix sort implementation from Intel's integrated performance primitives (IPP).

Additionally, a multi-threaded version of the FLiMS-based SIMD sort function is implemented, with the help of OpenMP pragmas. Both the merging function and the sort-in-chunks function are unaltered. The parallelisation is done on the calls of each, operating on equally-sized consecutive portions of the entire input, when possible.

The performance of the multi-threaded FLiMS-based sort is compared against the single-threaded baselines, as well as a parallel sort implementation in the Boost C++ libraries. The `block_indirect_sort()` function implements the samplesort sorting algorithm, and is regarded as one of the best performing C++ sort implementations [25].

Fig. 15 presents the results of both the single-threaded and multi-threaded experiments. The target processor is the 16-thread Ryzen Pro 4750U. The main observation is that the 16-thread FLiMS sort function surpasses the performance of the 16-thread `block_indirect_sort()` for the input range from 2^{17} to 2^{27} . A hybrid approach can be used to enable the single-threaded version of FLiMS for below 2^{20} to achieve the highest-performance overall, except radix sort.

Intel's radix sort takes the lead between 2^{12} and 2^{19} on this AMD processor, but IPP radix sort has some notable limitations. These include the less predictable performance across different systems when run across different processors (not demonstrated, for brevity), a restrictive license and radix-sort related implications, such as a limit of the input length to about 2^{28} for the tested implementation.

As a conclusion, FLiMS-based sorting is an attractive approach to accelerate merge sort using SIMD intrinsics.

However, it is still relevant to research solutions based on a variety of sorting methods, as different algorithms are more appropriate in different distributions and use cases [1]. For instance, radix sort can perform fewer data passes on data of a restricted range, such as for 10-bit integers.

9 FUTURE WORK

Current and future work includes developing efficient FLiMS implementations targeting other technologies such as GPUs, and exploring the adoption of FLiMS in various applications such as database analytics [2], [9] and beyond.

One concern in today's FPGA research on sorting is that it is mostly limited to fixed-width values [26]. Therefore, it would be helpful to also study FLiMS adaptations or alternatives for data of arbitrary width, such as strings [27].

The SIMD evaluation could also be extended with additional optimisation for specific processors and data. This could also include the skewness and stable variations and their applicability in this context. An AVX-512 version has already been developed, but for the target processor (Xeon 8124M) the performance benefits over AVX2 on the same processor were underwhelming. This possibly related to the efficiency of AVX-512 in the specific micro-architecture, though it would be appropriate to further investigate the applicability of FLiMS in future processor technologies.

It would also be interesting to formally prove any optimality FLiMS may exhibit, as well as to try to find equivalent circuits for merging more than two sorted input lists, for further reducing the size of parallel merge trees.

10 CONCLUSION

FLiMS is currently the most hardware-efficient 2-way merge block on FPGAs. It can be used to build efficient high-throughput merge trees for facilitating sorting of unsorted input on hardware. It features fewer and/or simpler pipeline stages than the alternatives, while achieving a higher amount of parallelism with less hardware resource utilisation. The skewness optimisation is the equivalent workaround found in an older merger, while maintaining the decentralised nature of the MAX entities logic, which is collectively used as a scalable selector logic. A variation for implementing stable sort is also presented to facilitate the needs of some database applications, as well as FLiMSj for unifying the dequeue signals that can be costly in some memory configurations. An SIMD implementation of FLiMS and its multi-threaded variant on a modern processor are also found to outperform popular highly-optimised C++ sort libraries. Sorting using such a high-throughput merge block can be more appropriate for big data applications than alternative approaches, since it yields streaming memory access patterns and can also be applied recursively for arbitrarily long data without keeping growing states.

REFERENCES

- [1] N. Auger, C. Nicaud, and C. Pivoteau, "Merge strategies: From merge sort to TimSort," *Working Paper or Preprint*, hal-01212839, Version 2, 2015. [Online]. Available: <https://hal-upec-upem.archives-ouvertes.fr/hal-01212839>

- [2] P. Papaphilippou, H. Pirk, and W. Luk, "Accelerating the merge phase of sort-merge join," in *Proc. 29th Int. Conf. Field Programmable Log. Appl.*, 2019, pp. 100–105.
- [3] W. Song, D. Koch, M. Luján, and J. Garside, "Parallel hardware merge sorter," in *Proc. 24th Annu. Int. Symp. Field-Programmable Custom Comput. Mach.*, 2016, pp. 95–102.
- [4] M. Saitoh, E. A. Elsayed, T. Van Chu, S. Mashimo, and K. Kise, "A high-performance and cost-effective hardware merge sorter without feedback datapath," in *Proc. IEEE 26th Annu. Int. Symp. Field-Programmable Custom Comput. Mach.*, 2018, pp. 197–204.
- [5] M. Saitoh and K. Kise, "Very massive hardware merge sorter," in *Proc. Int. Conf. Field-Programmable Technol.*, 2018, pp. 86–93.
- [6] E. A. Elsayed and K. Kise, "Towards an efficient hardware architecture for odd-even based merge sorter," in *Proc. IEEE 13th Int. Symp. Embedded Multicore/Many-Core Syst.-on-Chip*, 2019, pp. 249–256.
- [7] E. A. Elsayed and K. Kise, "High-performance and hardware-efficient odd-even based merge sorter," *IEICE Trans. Informat. Syst.*, vol. 103, no. 12, pp. 2504–2517, 2020.
- [8] P. Papaphilippou, C. Brooks, and W. Luk, "FLiMS: Fast lightweight merge sorter," in *Proc. Int. Conf. Field-Programmable Technol.*, 2018, pp. 78–85.
- [9] P. Papaphilippou, C. Brooks, and W. Luk, "An adaptable high-throughput FPGA merge sorter for accelerating database analytics," in *Proc. 30th Int. Conf. Field Programmable Log. Appl.*, 2020, pp. 65–72.
- [10] D. Koch and J. Torresen, "FPGASort: A high performance sorting architecture exploiting run-time reconfiguration on FPGAs for large problem sorting," in *Proc. 19th ACM/SIGDA Int. Symp. Field Programmable Gate Arrays*, 2011, pp. 45–54.
- [11] C.-Y. Lee and J.-M. Tsai, "A shift register architecture for high-speed data sorting," *J. VLSI Signal Process. Syst. Signal Image Video Technol.*, vol. 11, no. 3, pp. 273–280, 1995.
- [12] J. Casper and K. Olukotun, "Hardware acceleration of database operations," in *Proc. ACM/SIGDA Int. Symp. Field-Programmable Gate Arrays*, 2014, pp. 151–160.
- [13] R. Kobayashi and K. Kise, "Face: Fast and customizable sorting accelerator for heterogeneous many-core systems," in *Proc. IEEE 9th Int. Symp. Embedded Multicore/Many-Core Syst.-on-Chip*, 2015, pp. 49–56.
- [14] T. Usui, T. Van Chu, and K. Kise, "A cost-effective and scalable merge sorter tree on FPGAs," in *Proc. 4th Int. Symp. Comput. Netw.*, 2016, pp. 47–56.
- [15] K. Manev and D. Koch, "Large utility sorting on FPGAs," in *Proc. Int. Conf. Field-Programmable Technol.*, 2018, pp. 334–337.
- [16] K. E. Batcher, "Sorting networks and their applications," in *Proc. Spring Joint Comput. Conf.*, 1968, pp. 307–314.
- [17] J. Chhugani *et al.*, "Efficient implementation of sorting on multi-core SIMD CPU architecture," *VLDB Endowment*, vol. 1, no. 2, pp. 1313–1324, 2008.
- [18] A. Farmahini-Farahani, H. J. Duwe III, M. J. Schulte, and K. Compton, "Modular design of high-throughput, low-latency sorting units," *IEEE Trans. Comput.*, vol. 62, no. 7, pp. 1389–1402, Jul. 2013.
- [19] S. Mashimo, T. Van Chu, and K. Kise, "High-performance hardware merge sorter," in *Proc. 25th Annu. Int. Symp. Field-Programmable Custom Comput. Mach.*, 2017, pp. 1–8.
- [20] G. Zachmann, "Adaptive bitonic sorting," *Encyclopedia of Parallel Computing*, David Padua, Ed. Berlin, Germany: Springer, 2013, pp. 146–157.
- [21] The On-Line Encyclopedia of Integer Sequences, "Cullen numbers: $n * 2n + 1$," Accessed: Jun. 03, 2021. [Online]. Available: <https://oeis.org/A002064>
- [22] P. Papaphilippou, J. Meng, and W. Luk, "High-performance FPGA network switch architecture," in *Proc. Int. Symp. Field-Programmable Gate Arrays*, 2020, pp. 76–85.
- [23] B. Bramas, "A novel hybrid quicksort algorithm vectorized using AVX-512 on Intel Skylake," *Int. J. Adv. Comput. Sci. Appl.*, vol. 8, no. 10, pp. 337–344, 2017.
- [24] Intel (R), Intel intrinsics guide. Accessed: Nov. 10, 2021. [Online]. Available: <https://software.intel.com/sites/landingpage/IntrinsicsGuide/>
- [25] "Boost.sort 3.- parallel algorithms," Accessed: Apr. 09, 2021. [Online]. Available: <https://www.boost.org/doc/libs/develop/libs/sort/doc/html/sort/parallel.html>
- [26] T. Van Chu, "FPGA dramatically speeds up sorting of basic algorithms (5)," *ACRI*, Tokyo Inst. Technol. Accessed: Nov. 10, 2021. [Online]. Available: <https://www.acri.c.titech.ac.jp/wordpress/archives/4713>
- [27] M. Asiatici, D. Maiorano, and P. Ienne, "How many CPU cores is an FPGA worth? Lessons learned from accelerating string sorting on a CPU-FPGA system," *J. Signal Process. Syst.*, vol. 93, pp. 1–13, 2021.



Philippos Papaphilippou received the PhD degree from Imperial College London. His PhD was funded by dunhumby for researching novel accelerators to improve the performance of big data analytics. His research interests include FPGAs, sorting algorithms, network switches, multi-processor architecture and data science.



Wayne Luk (Fellow, IEEE) is currently a professor of computer engineering with Imperial College London. He leads the Programming Languages and Systems Section, and the Custom Computing Research Group with the Department of Computing. He is a fellow of the Royal Academy of Engineering and the BCS.



Chris Brooks is currently the head of Science Innovation with dunhumby, U.K. He is accountable for researching and developing new science techniques and technical implementations. He leads science research into a variety of areas, focussed primarily on the retail domain.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.